

***Astropulse: A Search for Microsecond Transient Radio Signals
Using Distributed Computing. I. Methodology***

J. Von Korff^{1 2 3}, *P. Demorest*⁴, *E. Heien*^{1 5}, *E. Korpela*¹, *D. Werthimer*¹, *J. Cobb*¹,
*M. Lebofsky*¹, *D. Anderson*¹, *B. Bankay*¹, *A. Siemion*¹

Received _____; *accepted* _____

¹Space Sciences Lab, University of California, Berkeley, CA 94720

²Physics department, Kansas State University, Manhattan, KS 66502

³vonkorff@gmail.com

⁴National Radio Astronomy Observatory, Charlottesville, VA 22903

⁵Graduate School of Information Science and Technology, Osaka University

ABSTRACT

We are performing a transient, microsecond timescale radio sky survey, called “Astropulse,” using the Arecibo telescope. Astropulse searches for brief ($0.4\ \mu\text{s}$ to $204.8\ \mu\text{s}$), wideband (relative to its 2.5 MHz bandwidth) radio pulses centered at 1,420 MHz. Astropulse is a commensal (piggyback) survey, and scans the sky between declinations of -1.33 and 38.03 degrees. We obtained 1,540 hours of data in each of 7 beams of the ALFA receiver, with 2 polarizations per beam. The data are 1-bit complex sampled at the Nyquist limit of $0.4\ \mu\text{s}$ per sample. Examination of timescales on the order of a few microseconds is possible because we used coherent dedispersion, a technique that has frequently been used for targeted observations, but has never before been associated with a radio sky survey. The more usual technique, incoherent dedispersion, cannot resolve signals below a minimum timescale which depends on the dispersion measure and frequency of the signal. However, coherent dedispersion requires more intensive computation than incoherent dedispersion. The required processing power was provided by BOINC, the Berkeley Open Infrastructure for Network Computing. BOINC is a distributed computing system, which allows us to utilize hundreds of thousands of volunteers’ computers to perform the necessary calculations for coherent dedispersion. Astrophysical events that might produce brief radio pulses include giant pulses from pulsars, RRATs, exploding primordial black holes, or new sources yet to be imagined. Radio frequency interference (RFI) and noise contaminate the data; these are mitigated by a number of techniques including multi-polarization correlation, DM repetition detection, and frequency profiling.

Subject headings: radio continuum: general — extraterrestrial intelligence — pulsars: general — black hole physics — cosmology: early universe

Facility: Arecibo(ALFA)

1. Introduction

1.1. Scientific motivation

This is an exciting time in the field of transient astronomy, both in the radio and in other parts of the spectrum. Improving technology allows astronomers to perform fast followups of transient events, store extensive digital records of observations, and run processor-intensive algorithms on data in real time. These advances make possible instruments that examine optical afterglows of gamma-ray bursts (Vestrand et al. 2005) or neutrino sources (Kowalski & Mohr 2007). High resolution digital images can be recorded and stored quickly using current (Kaiser 2004) and planned technology (Ivezic et al. 2008). In the radio, astronomers search for transients such as orphan GRB afterglows (Levinson et al. 2002) or radio bursts of unknown origin (Katz et al. 2003).

Our project, called “Astropulse,” searches for brief, wideband radio pulses on timescales of microseconds to milliseconds, and surveys the entire sky visible from Arecibo Observatory. The idea of a short-timescale radio observation is not new. Other experiments are well-suited for detecting radio pulses on a microsecond timescale, or even much shorter scales. However, these observations are directed; they examine known phenomena. For instance, such an experiment might record the nanosecond structure of the signals from the Crab pulsar. And of course the idea of a radio survey is not new. Other experiments perform surveys for radio pulses over large regions of the sky. However, these observations examine $50\ \mu\text{s}$ timescales or longer. Astropulse is the first radio survey for transient phenomena with microsecond resolution.

This project is made possible by Astropulse’s access to unprecedented processing power, using the distributed computing technique. Because the interstellar medium disperses radio signals, all of our data must be dedispersed. We send our data to volunteers, who perform coherent dedispersion (Lorimer & Kramer 2005; Hankins et al. 1987) using their own computers. Then they send the results of this computation back to us, informing us whether they detected a signal, and reporting that signal’s dispersion measure, power, and other parameters. Astropulse is processor intensive because we must perform coherent dedispersion, whereas other surveys perform incoherent dedispersion. Coherent dedispersion is necessary to resolve structures below $50\ \mu\text{s}$ or so, depending on the dispersion measure.

We are not committed to detecting any particular astrophysical source; rather, we are motivated by our ability to examine an unexplored region of parameter space. However, we

consider that we might detect evaporating primordial black holes, millisecond (or faster) pulsars, or RRATs (rotating radio transients, McLaughlin et al. 2006). We will consider each of these possibilities in turn. We could potentially detect pulsed communications from extraterrestrial civilizations, though we do not discuss this possibility herein.

1.2. Black holes

1.2.1. Hawking radiation

It was proposed by Hawking (1974) that a black hole of mass M emits radiation like a black body whose temperature is given by the following relation:

$$T_{BH} = \frac{\hbar c^3}{8\pi k G M} = 10^{-6} \left(\frac{M_{\odot}}{M} \right) \text{ K}. \quad (1)$$

The radiant energy comes directly from the black hole’s mass, and as a result, it is losing mass at a rate $\dot{M} \propto -M^{-2}$. Because the black hole radiates more power as it shrinks, we expect a burst of energy in the last moments of the black hole’s life. One can make different assumptions about the energy distribution of the radiation from a black hole evaporation (Carter et al. 1976). For a “hard” equation of state, with an adiabatic index $\Gamma > \frac{6}{5}$, the radiation does not reach thermal equilibrium. The standard model falls into this category, and it would assume that the radiation behaves as a relativistic ideal gas, $\Gamma = \frac{4}{3}$. In this case, the final explosion of the black hole lasts on the order of seconds. However, for a “soft” equation of state, as proposed by Hagedorn (1965), Γ could be much smaller. In this case, the explosion might happen in 10^{-7} seconds or less. Astropulse is ideally suited for detecting such fast explosions.

We can integrate the radiant energy to find the total lifetime of the black hole, demonstrating that if the black hole is exploding now, it must have been created at a mass of 10^{12} kg or less. Such a small black hole cannot have been born from a star, and would have to be created in the big bang (Hawking 1971), from “density perturbations in the early universe” (MacGibbon et al. 1990).

1.2.2. *Electromagnetic pulses*

The total amount of energy released in the last second of the black hole’s life is about 10^{23} J. Most previous studies have attempted to detect this energy in the cosmic gamma ray background (Raine & Thomas 2005). But Rees (1977) suggested that some of this energy could be converted into a radio pulse. The idea is that as the black hole shrinks and becomes hotter, it starts radiating more and more massive particles, including electrons and positrons (due to pair production at the event horizon), but later, heavier particles as well. This forms a plasma fireball expanding around the black hole. As this conducting shell expands into the ambient magnetic field, it pushes the field out of the way, creating an electromagnetic pulse. Rees argued that for a magnetic field B around 5×10^{-6} Gauss and a critical mass of $\sim 2 \times 10^{11}$ g, a radio pulse detectable in the 21 cm band is plausible.

An observation of these pulses would be a very significant confirmation of both Hawking radiation and the existence of primordial black holes (PBHs). At the very least, we can put a limit on the possible maximum density of evaporating black holes in the universe, if we make some assumptions about their distribution, and contingent on the assumption that they produce radio pulses. This information would be relevant to cosmological models describing the big bang.

Some groups have searched for these PBHs in the radio, but many researchers have looked for gamma-ray emission instead (Ukwatta et al. 2010). Radio and gamma-ray surveys make very different assumptions about the PBHs’ evaporation time, so their results are difficult or impossible to compare in a meaningful way. Ukwatta et al. describe PBH explosions as having timescales of seconds or minutes, whereas Astropulse is looking for microsecond pulses.

1.3. Other sources

Astropulse might also detect RRATs (McLaughlin et al. 2006) or repeating or giant pulses from pulsars. Of these possibilities, giant pulses are the most likely. Astropulse is optimized for pulses of $200 \mu s$ or less, but its sensitivity relative to other surveys is best at short timescales, around 0.4 to $1.6 \mu s$. These timescales are much too short for repeating pulses from a pulsar (Lattimer et al. 1990), even a millisecond pulsar with a small beam opening angle half-width, or from a RRAT.

Giant pulses, on the other hand, can have very short timescales suitable for detection by Astropulse. The Crab’s giant pulses have structure ranging from a few nanoseconds in duration at $2 \text{ Jy } \mu s$ (Hankins et al. 2003) to $64 \mu s$ or more at 1,000 to 10,000 Jy μs , with

a typical pulse duration of a few microseconds (Popov & Stappers 2007). (See Section 3.1 for a discussion of the $\text{Jy } \mu\text{s}$ unit.)

2. Telescope and instrumentation

2.1. Sky coverage

Arecibo Observatory scans approximately one third of the sky, between declinations of -1.33 and 38.03 degrees. Because of this, Astropulse cannot see the galactic center (around -29° dec) but can see 452 out of the 1826 pulsars in the ATNF pulsar database¹, including the Crab. Astropulse is a commensal survey; this means that other surveys control the telescope pointing, but allow Astropulse to collect data at all times. Our partner surveys include GALFA (Galactic ALFA), discussed in Peek & Heiles (2008); Stanimirovic & Putnam (2006) and PALFA (Pulsar ALFA), discussed in Cordes (2008). Our group also operates SETI@home, another commensal radio survey, and the two projects use the same data: a 1-bit complex sampled 2.5 MHz bandwidth centered at 1420 MHz. To date, we have observed for 1,540 hours with each of the 7 beams (and 2 linear polarizations per beam), for a total of 21,600 hours of observation time. We have been taking our primary set of data using the ALFA receiver from September 2006 until May 2010, for a total of 3.7 years. This implies that we have had 1/21 of all possible observation time during those years. Since a good deal of Arecibo’s time is dedicated to non-astronomical purposes, such as ionospheric science, our fraction of astronomy time is significantly larger than 1/21.

2.2. ALFA receiver

The ALFA (Arecibo L-band Feed Array) receiver has 7 dual-polarization beams on the sky arranged in a hexagonal pattern, each with a $3.5'$ beamwidth. The central beam has a gain of 11 K / Jy, and the other beams have 8.5 K / Jy². The system temperature is 30 K. The 6 peripheral beam pointings differ from the central beam by a maximum of $6.4'$.

¹<http://www.atnf.csiro.au/research/pulsar/psrcat/>, as of 6/29/2009

²http://www.naic.edu/alfa/gen_info/info_obs.shtml

2.3. Downconverter

Multiple experiments use the signal from the ALFA receiver, so we split the signal using an IF splitter. These 14 signals are attenuated by 6 to 13 decibels for purposes of level-matching, and then enter our multibeam quadrature baseband downconverter. Downconversion involves complex multiplication, resulting in 14 complex or 28 real channels.

2.4. Data recorder

These 28 real channels are digitized with 1 bit precision using comparators, and the resulting digital signals (the signs of the 28 voltages) are directed through ribbon cables to Digital Data Acquisition (DDA) cards on a PC, which is running our software that acquires and writes the data to disk. In addition, this software collects telescope coordinates from the Arecibo telescope’s data broadcast network, SCRAMnet. The coordinates consist of the Right Ascension (RA) and Declination (Dec) to which the telescope is currently pointing, as well as the time for which that RA and Dec are valid.

The data files are stored on a hot swappable SATA drive, which fills up in 14 to 20 hours of observation time. Since we are taking data 1/21 of the time, we must swap out the SATA drive about once per two weeks. When enough drives have been collected, the Arecibo staff ship the drives to us at Space Sciences Lab, UC Berkeley. We use 20 SATA drives in all, each of which holds 500 or 750 GB. We also send a backup copy of each file to NERSC, the National Energy Research Scientific Computing Center. This ensures that we can retrieve the data at any time. In all, we have taken over 48 TB of data from ALFA multibeam. We need a large (6 TB) disk array at Berkeley to buffer the data before it is sent to volunteers. The volunteers’ PCs then process the data and send the results back to Berkeley. For a discussion of data processing after this point, see Section 4 on BOINC, and Section 3 on the dedispersion algorithm for the volunteers’ client program.

3. Pulse detection : thresholds and dedispersion

3.1. Overview

The primary function of the Astropulse program is to dedisperse potential pulsed signals (“candidate pulses”), then determine whether the dedispersed pulse surpasses an

appropriate power threshold. We will discuss the theory behind dedispersion and then the methods we use to select the thresholds. Then we will calculate the sensitivity of Astropulse in $\text{Jy } \mu\text{s}$, a unit of “pulse area.”

The $\text{Jy } \mu\text{s}$ unit refers to the pulse’s flux density (in Jy) integrated over its duration (in μs .) It is called an “area” because it is calculated using this integral, which is the area under a curve. Although the unit of flux density (Jy) is a more conventional measure of sensitivity, the $\text{Jy } \mu\text{s}$ is more meaningful in our case, because we are attempting to detect unresolved pulses. For example, consider two pulses; one is 500 Jy and lasts $0.2 \mu\text{s}$, and the other is 1,000 Jy and lasts $0.1 \mu\text{s}$. When these pulses are dispersed, they will be similar in appearance; Astropulse cannot distinguish between them because their dedispersed durations are shorter than Astropulse’s time resolution. But Astropulse can determine that both pulses are $100 \text{ Jy } \mu\text{s}$.

Note that when we describe a measured pulse’s apparent area in $\text{Jy } \mu\text{s}$, the actual pulse area may be different depending on any contributions to the system temperature. We assume a particular minimal system temperature (30 K) for the ALFA receiver, whereas we might have a different effective system temperature when looking at the Crab nebula.

3.2. Dedispersion

Between a radio pulse’s source (i.e. black hole, pulsar, or ET) and our detector, the pulse is dispersed as it travels through the Interstellar Medium (ISM). According to Lorimer & Kramer (2005), the relative time delay for frequency ν is given by:

$$t(\nu) = \mathcal{D} \times \text{DM} / \nu^2 \tag{2}$$

$$\text{DM} = \int n_e d\ell \tag{3}$$

where d is the distance to the source of the pulse, n_e is the electron density, and \mathcal{D} is equal to $4.15 \times 10^3 \text{ MHz}^2 \text{ pc}^{-1} \text{ cm}^3 \text{ s}$.

A useful estimate for the dispersion measure weighted mean electron density in our Galaxy is $n_e = 0.03 \text{ cm}^{-3}$ (Guélin 1973).

Astropulse loops through the data at several nested levels, and considers DMs ranging

from -830 pc cm^{-3} to -49.5 pc cm^{-3} and from 49.5 pc cm^{-3} to 830 pc cm^{-3} . We chose the lower limit, 49.5 pc cm^{-3} , for two reasons. First, we found that our sensitivity diminishes at low dispersion measures due to the effects of one-bit digitization. For instance, a (hypothetical) very strong, undispersed, $2 \mu\text{s}$ signal would be undetectable, since its signature in our data would only be five samples long, and each sample carries only one bit of information. Therefore, we can only detect dispersed pulses. Second, local interference at Arecibo Observatory is stronger at low dispersion measures. We tested the first effect by inserting simulated pulses into our detection algorithm, and the second effect by examining data from the telescope. In this way, we empirically determined a lower limit for our dispersion measure. The upper limit of 830 pc cm^{-3} was selected so that approximately half of the volume of the Galactic plane would be visible to our search according to the Galactic maps in Cordes & Lazio (2003).

Astropulse considers pulses of widths ranging from $0.4 \mu\text{s}$ (a single sample) to $204.8 \mu\text{s}$. The larger widths are tested by summing 2^ℓ adjacent samples after dedispersion, where $0 \leq \ell \leq 9$, and ℓ takes integer values.

3.2.1. Incoherent dedispersion and its limitations

We have two choices for our methodology: coherent dedispersion and incoherent dedispersion. Astropulse uses coherent dedispersion, whereas other radio surveys use incoherent dedispersion. Incoherent dedispersion is much more computationally efficient, and for longer timescales it’s almost as good as coherent dedispersion. However, as we will see, Astropulse would be unable to examine the $0.4 \mu\text{s}$ timescale without coherent dedispersion.

Incoherent dedispersion means that the signal’s power spectrum is calculated, and the power vs. time of each sub-band is analyzed. The method is called “incoherent” for this reason – the phase information about individual frequencies is lost; only the total power of each sub-band at each time is retained. Next, the sub-bands are realigned at all possible dispersion measures, in an effort to find one DM at which the components align to produce a large power in a short period of time.

However, incoherent dedispersion is limited in two ways. First, the goal of recording power vs. time makes sense only on a timescale greater than

$$dt_1 = \frac{1}{d\nu}, \quad (4)$$

where $d\nu$ is the width of each sub-band. This is because of time-frequency uncertainty.

Second, in each sub-band the pulse is dispersed by some amount $dt_2(\nu)$, which we can find using Equation 2.

$$dt_2(\nu) = t(\nu) - t(\nu + d\nu) = \frac{\mathcal{D} \cdot \text{DM}}{(\nu + d\nu)^2} - \frac{\mathcal{D} \cdot \text{DM}}{\nu^2} \quad (5)$$

$$= \frac{\mathcal{D} \cdot \text{DM}}{\nu^2} \left(\frac{\nu^2}{(\nu + d\nu)^2} - 1 \right) \quad (6)$$

$$= \frac{\mathcal{D} \cdot \text{DM}}{\nu^2} \left(\left(\frac{\nu}{\nu + d\nu} \right)^2 - 1 \right) \quad (7)$$

$$= \frac{\mathcal{D} \cdot \text{DM}}{\nu^2} \left(\left(1 - \frac{d\nu}{\nu + d\nu} \right)^2 - 1 \right) \quad (8)$$

Under the assumption that incoherent dedispersion divides the band up into many small pieces, $d\nu \ll \Delta\nu$, or that the bandwidth is much smaller than the frequency, $\Delta\nu \ll \nu$, we obtain $d\nu \ll \nu$. Then Equation 8 becomes:

$$dt_2(\nu) \approx \frac{\mathcal{D} \cdot \text{DM}}{\nu^2} \left(1 - 1 + 2\frac{d\nu}{\nu} \right) \quad (9)$$

$$= 2(\mathcal{D} \cdot \text{DM}) \frac{d\nu}{\nu^3} \quad (10)$$

Because the signal in each sub-band has a differential dispersion delay of $dt_2(\nu)$, the method cannot localize the pulse better than this. (If the survey examines a large bandwidth $\Delta\nu$, $dt_2(\nu)$ may vary substantially for different sub-bands, but we are considering each sub-band individually.) For a signal of zero width, the differential dispersion delay across the band would be equal to the duration of the dispersed pulse. However, the component of the signal in each sub-band does not have zero width – it has width dt_1 . We will assume that the minimum timescale for incoherent dedispersion, dt , cannot be less than either dt_1 or $dt_2(\nu)$, so that the smallest possible width is for $dt = dt_1 = dt_2(\nu)$. (This assumption is appropriate if the signal has a slowly-varying power as a function of frequency, and if all frequencies were emitted at the same time prior to the dispersion process in the ISM. See Section 3.2.3 for a related argument.) Then we find that the minimum timescale for incoherent dedispersion happens when:

$$2(\mathcal{D} \cdot \text{DM}) \frac{d\nu}{\nu^3} = \frac{1}{d\nu} \quad (11)$$

$$d\nu^2 = \frac{\nu^3}{2\mathcal{D} \cdot \text{DM}} \quad (12)$$

$$d\nu = \sqrt{\frac{\nu^3}{2\mathcal{D} \cdot \text{DM}}} \quad (13)$$

$$dt(\nu) = \sqrt{\frac{2\mathcal{D} \cdot \text{DM}}{\nu^3}} \quad (14)$$

where the last step used Equation 4. For the Crab pulsar, the dispersion measure is 56.8 pc cm^{-3} (Sallmen et al. 1999), and we are observing at a frequency of 1.42 GHz. Then we can substitute $\mathcal{D} = 4.15 \times 10^3 \text{ MHz}^2 (\text{pc cm}^{-3})^{-1} \text{ s}$, $\text{DM} = 56.8 \text{ pc cm}^{-3} \times \frac{DM}{56.8 \text{ pc cm}^{-3}}$, and $\frac{1}{\nu} = \frac{1420 \text{ MHz}}{\nu} \times \frac{1}{1420 \text{ MHz}}$ to obtain:

$$dt(\nu) = \sqrt{(8.3 \times 10^3 \text{ MHz}^2 (\text{pc cm}^{-3})^{-1} \text{ s}) \times 56.8 \text{ pc cm}^{-3}} \quad (15)$$

$$\times \sqrt{\frac{\text{DM}}{56.8 \text{ pc cm}^{-3}} \left(\frac{1420 \text{ MHz}}{\nu}\right)^3 \times \left(\frac{1}{1420 \text{ MHz}}\right)^3} \quad (16)$$

$$= 12.8 \text{ } \mu\text{s} \left(\frac{DM}{56.8 \text{ pc cm}^{-3}}\right)^{0.5} \left(\frac{\nu}{1.42 \text{ GHz}}\right)^{-1.5}. \quad (17)$$

So for the Crab pulsar, this is a limit of $12.8 \text{ } \mu\text{s}$, or 32 samples at 1.42 GHz. For a more distant source, the limit might be as much as $50 \text{ } \mu\text{s}$, or 124 samples. (Astropulse considers sources with a DM as high as 830 pc cm^{-3} .)

3.2.2. Coherent dedispersion as deconvolution

Coherent dedispersion (Lorimer & Kramer 2005; Hankins et al. 1987) is an alternative technique that allows better time resolution by performing the mathematical inverse of the ISM's dispersion operation. Coherent dedispersion deals with amplitude rather than power, preserving phase information. In the absence of noise or scattering, and given precise knowledge of the pulse's dispersion measure, coherent dedispersion would reconstruct the

original pulse exactly. We need to analyze the mathematical operation corresponding to dispersion in order to find its inverse.

If $F(n)$ is the original pulse as a function of sample number n , suppose $D[F]$ is the dispersed pulse. Then, relying on the time translation invariance and linearity of the dispersion operator, we can show that D is just a convolution. In particular, $D[F] = F * D[\delta]$, where $*$ is convolution and δ is the discrete δ function. Finally, the convolution theorem for Fourier transforms gives us:

$$F = \text{DFT}^{-1}\left(\frac{\text{DFT}(D[F])}{\text{DFT}(D[\delta])}\right). \quad (18)$$

Equation 18 gives us a fast method for dedispersing a pulse $D[F]$, obtaining the original pulse F . This method is fast because the Fast Fourier Transform (FFT) algorithm is fast, taking time $O(N \log N)$ to Fourier transform N samples of data. We can use this fact to estimate the run time of Astropulse’s dedispersion algorithm. Astropulse has to dedisperse each set of N samples many times, since one dedispersion must be performed for every dispersion measure. If M is the number of dispersion measures to be tested, then N samples can be dedispersed in time $O(MN \log N)$. Furthermore, Astropulse must operate on a long stream of data, of length L , which is much longer than the length N of a single Fourier transform. Therefore, the total time required is $O(ML \log N)$.

A similar calculation suggests that incoherent dedispersion would be faster. This is partly because incoherent dedispersion performs fewer tests. Incoherent dedispersion cannot test as many dispersion measures as coherent dedispersion does, because its time resolution is imperfect and it cannot always distinguish between different dispersion measures. To test a particular dispersion measure, an algorithm must target a particular time delay between the minimum and maximum frequency in the band. This time delay cannot be determined more accurately than the time resolution dt (Equation 17) resulting from incoherent dedispersion. If dt corresponds to n samples, then incoherent dedispersion can test only $\frac{1}{n}$ as many dispersion measures as coherent dedispersion. In other words, with incoherent dedispersion, we would test not M , but M/n dispersion measures.

With incoherent dedispersion, we must process L samples for each dispersion measure, so we require time $O(ML/n)$. (Some additional time is required to Fourier transform the data into a power spectrum, but this process is not dominant.) So the time ratio between coherent and incoherent dedispersion is $O(ML \log N / (ML/n)) = O(n \log N)$. In our case, the time resolution dt will vary by dispersion measure, but assuming a value of $20 \mu\text{s}$, corresponding to a DM of 139 pc cm^{-3} , and a sample duration of $0.4 \mu\text{s}$, we obtain $n = 50$. Then $\log N = 15$, so $n \log N = 750$, a large number. Therefore, we expect that incoherent dedispersion is substantially faster than coherent dedispersion, as long as both methods are

applicable. Coherent dedispersion is useful in situations where a very short time resolution is required, shorter than that permitted by incoherent dedispersion.

3.2.3. *Computing the nonlinear chirp function*

The “chirp” function, $\text{DFT}(D[\delta])$, plays an important role in Equation 18. The function $D[\delta]$ represents a dispersed delta function, and the DFT transforms it into the frequency domain. In other words, we imagine a brief, strong pulse emitted by an astrophysical source. The pulse is dispersed by its passage through the interstellar medium, resulting in a chirp function that is spread out in time. (The name “chirp” is meant to suggest a sound with changing frequency, just as the chirp signal has a changing radio frequency.) Our goal in this section will be to compute the functional form of this dispersed signal in the frequency domain.

For the sake of brevity, we will write f for $D[\delta]$, and \tilde{f} for its Fourier transform, $\text{DFT}(D[\delta])$. Then $f(t)$ can be written in terms of an amplitude and a phase:

$$f(t) = A(t) \exp i\theta(t) \tag{19}$$

If we assume that $A(t)$ is slowly-varying, then it is meaningful to talk about a “frequency at time t ”, $\omega(t) = \frac{d\theta}{dt}$, or $\nu(t) = \frac{d\theta}{dt}/2\pi$. In that case, we could compute $\theta(t) = 2\pi \int \nu(t) dt$. Analogously, we will assume that $\tilde{f}(t)$ can be written in terms of an amplitude and a phase:

$$\tilde{f}(\nu) = \tilde{A}(\nu) \exp i\tilde{\theta}(\nu) \tag{20}$$

Using similar logic, we could hypothesize that $\tilde{\theta}(\nu) = 2\pi \int t(\nu) d\nu$. This hypothesis turns out to be correct, except for a minus sign, which we will now demonstrate.

We begin by finding the Taylor expansion of $\tilde{\theta}(\nu)$ around a particular frequency ν_1 . This gives us the expression:

$$\tilde{\theta}(\nu) \approx \tilde{\theta}(\nu_1) + \left. \frac{d\tilde{\theta}}{d\nu} \right|_{\nu_1} \cdot (\nu - \nu_1) \quad (21)$$

Equation 21 corresponds to the component of the signal with frequencies near ν_1 . We can determine the arrival time of these frequencies by finding the Fourier transform of the frequency distribution:

$$g(t) = \int \tilde{A}(\nu) \exp(i(\tilde{\theta}(\nu_1) + \left. \frac{d\tilde{\theta}}{d\nu} \right|_{\nu_1} \cdot (\nu - \nu_1))) \exp(i2\pi\nu t) d\nu \quad (22)$$

$$= \int \tilde{A}(\nu) \exp(i(\left. \frac{d\tilde{\theta}}{d\nu} \right|_{\nu_1} + 2\pi t)\nu) \exp(-i \left. \frac{d\tilde{\theta}}{d\nu} \right|_{\nu_1} \cdot \nu_1) d\nu \quad (23)$$

When $t = -\frac{1}{2\pi} \left. \frac{d\tilde{\theta}}{d\nu} \right|_{\nu_1}$, the integrand of Equation 23 becomes constant, so the integral is infinite. This, then, is $t(\nu_1)$, the arrival time of frequency ν_1 . Therefore,

$$\left. \frac{d\tilde{\theta}}{d\nu} \right|_{\nu_1} = -2\pi t(\nu_1) \quad (24)$$

$$\tilde{\theta}(\nu) = -2\pi \int t(\nu) d\nu \quad (25)$$

This demonstrates that our hypothesis was correct, except for a minus sign. So to compute the phase of $\tilde{f}(\nu)$, the chirp function in the frequency domain, we need only find $t(\nu)$ and integrate. We can rewrite Equation 2 by combining the constant factors:

$$t(\nu) = B/\nu^2 \quad (26)$$

Since this equation gives the relative time delay rather than the absolute time delay, it will still be true if we add or subtract a constant of our choice. Using band center frequency $\nu_0 = 1420$ MHz, we obtain:

$$t(\nu) = B\left(\frac{1}{\nu^2} - \frac{1}{\nu_0^2}\right). \quad (27)$$

We can find the exponent of the frequency-domain amplitude by integrating this equation. After including the factor of -2π from Equation 25, we obtain:

$$-2\pi \int t(\nu) = 2\pi B\left(\frac{1}{\nu} + \frac{\nu}{\nu_0^2} + C\right) \quad (28)$$

$$= 2\pi B\left(\frac{\nu_0^2}{\nu\nu_0^2} + \frac{\nu^2}{\nu\nu_0^2} + C\frac{\nu_0^2\nu}{\nu\nu_0^2}\right) \quad (29)$$

$$= 2\pi B\frac{(\nu - \nu_0)^2}{\nu\nu_0^2}. \quad (30)$$

where we have judiciously chosen $C = -2$ to simplify the equation. This results in a frequency domain amplitude:

$$\tilde{f}(\nu) = \tilde{A}(\nu) \exp\left(2\pi i B \frac{(\nu - \nu_0)^2}{\nu\nu_0^2}\right). \quad (31)$$

Since $\nu \approx \nu_0$ in our application, the chirp function is approximately $\exp(2\pi i B \frac{(\nu - \nu_0)^2}{\nu_0^3})$, so that the exponent can be approximated as a quadratic in ν . But the extra factor ν_0/ν is easily included in our dedispersion algorithm, so we have done so. This extra factor is close to 1, since in our case $\nu_0 = 1.42$ GHz, and ν differs from ν_0 by at most 1.25 MHz. (We can ignore the factor $\tilde{A}(\nu)$, assuming it is approximately constant in ν over our small bandwidth.)

3.3. Thresholds

Astropulse searches for pulses whose power exceeds certain thresholds. These thresholds can be calculated either experimentally or theoretically. We'll start by finding the theoretical values, then point out some of the uncontrollable factors that make these values inaccurate, and finally we'll describe a Monte Carlo simulation method for calculating thresholds.

3.3.1. Single pulse thresholds: theory

We want to calculate the distribution (pdf) of the integrated noise power in 2^ℓ samples, after dedispersion. (Here, the “power” refers to the absolute value of the square of the amplitude, where the undispersed time series has amplitudes ± 1 .) We will perform some calculations and conclude that the dedispersed signal $f_d(t)$ is distributed like a complex Gaussian at each time t .

First, we assume that the pre-dedispersion time series is pure white noise; that is, each bit of a two bit complex sample is independently distributed with equal probability of a 0 or 1, so each $f(t)$ has equal probability for $\pm 1 \pm i$. (Note that this f is discrete-valued with a discrete time argument, as opposed to the continuous-valued f with continuous time argument described in Section 3.2.3.) Then we deconvolve this data by FFT. In other words,

$$\tilde{f}(k) = \frac{1}{\sqrt{N}} \sum_{t=0}^{N-1} f(t) e^{-2\pi i k t / N}. \quad (32)$$

The distribution of a single $\tilde{f}(k)$ is Gaussian (by the central limit theorem), and to deduce its variance, we will find the variance of its real and complex components $\Re(\tilde{f}(k))$ and $\Im(\tilde{f}(k))$ independently:

$$\Re(\tilde{f}(k)) = \frac{1}{\sqrt{N}} \sum_{t=0}^{N-1} \Re(f(t)) \cos(2\pi k t / N) - \Im(f(t)) \sin(2\pi k t / N) \quad (33)$$

$$\begin{aligned} \text{Var}(\Re(\tilde{f}(k))) &= \left(\frac{1}{\sqrt{N}}\right)^2 \sum_{t=0}^{N-1} \left(\text{Var}(\Re(f(t))) \cos^2(2\pi k t / N) \right. \\ &\quad \left. - \text{Var}(\Im(f(t))) \sin^2(2\pi k t / N) \right) \end{aligned} \quad (34)$$

$$= \frac{1}{N} \sum_{t=0}^{N-1} \cos^2(2\pi k t / N) + \sin^2(2\pi k t / N) \quad (35)$$

$$= \frac{1}{N} \cdot N = 1 \quad (36)$$

Equation 35 follows from the previous equation because $f(t)$ can only take the values

$\pm 1 \pm i$, therefore its real and imaginary components each have variance of 1.

Therefore, the variance of the real component of $\tilde{f}(k)$ is 1, and the same argument holds for the imaginary component. Once we have obtained $\tilde{f}(k)$, the remaining steps in the dedispersion are to multiply by a frequency-domain chirp function, followed by Fourier transforming back to the time domain.

The frequency-domain chirp function has the form $e^{i\tilde{\theta}(k)}$ for some real phase $\tilde{\theta}$. Since $\tilde{f}(k)$ is already a complex number with random phase, multiplication by another complex phase has no effect on the probability distribution of $\tilde{f}(k)$. Finally, we run the inverse Fourier transform to obtain a dedispersed signal, $f_d(t)$. Since $\text{Var}(\Re(\tilde{f}(k))) = \text{Var}(\Re(f(t))) = 1$, the same mathematical argument that we applied to the forward Fourier transform also applies to the inverse Fourier transform. Then we find that the dedispersed signal $f_d(t)$ is again a complex Gaussian, such that $\Re(f_d(t))$ and $\Im(f_d(t))$ are Gaussians with variance 1.

The power in each sample after dedispersion will be distributed as $|\Re(f_d(t))|^2 + |\Im(f_d(t))|^2$, the sum of the squares of two standard Gaussians. This distribution is easily calculated; the joint probability distribution is:

$$\frac{1}{\sqrt{2\pi}}e^{-\frac{x^2}{2}} \cdot \frac{1}{\sqrt{2\pi}}e^{-\frac{y^2}{2}} dx dy \quad (37)$$

$$= \frac{1}{2\pi}e^{-\frac{r^2}{2}} r dr d\theta \quad (38)$$

$$\rightarrow e^{-\frac{r^2}{2}} r dr \quad (39)$$

$$= e^{-u} du. \quad (40)$$

where $u = \frac{x^2+y^2}{2}$ is half the power in one sample, in the time domain. Therefore, half the power is exponentially distributed with mean 1; or equivalently, the power is exponentially distributed with mean 2.

In future calculations, we will normalize to half of this power, so that the average power per sample is 1.

So for instance, if after dedispersion we find that a certain sample has a power P , we conclude that only one in e^P samples has a comparable power. To ascertain how unlikely this is, we need to calculate how many such samples we have examined over the entire course of the experiment. This would be:

$$\begin{aligned}
& 48 \text{ TB} \times 4 \cdot 10^{12} \text{ samples per TB} \times 14208 \text{ DMs} \\
& \quad \times 2 \text{ DM signs} \\
& = 5.45 \times 10^{18} = e^{43.1}.
\end{aligned} \tag{41}$$

So far, we have considered pulses that are one sample in width. However, we are also searching for pulses of width 2, 4, 8, \dots , 512 samples. To search for wider pulses, we sum the power over 2^ℓ adjacent samples, for each integer ℓ between 0 and 9. We will refer to this summation as “co-adding,” and each of the 10 possible widths is a “co-add.” We compare this summed power to an appropriate threshold, which is larger at higher co-adds. There are half as many potential pulses at each co-add, compared with the previous co-add. So Equation 41 underestimates the number of potential pulses by a factor of $1 + \frac{1}{2} + \frac{1}{4} + \dots + \frac{1}{512} \approx 2$.)

Therefore, with a threshold of 43.8 for one-sample potential pulses, and appropriate thresholds for potential pulses of other sizes, we would rule out all but one noise event over the course of our entire observation history. Thus, if Astropulse had no mechanism to remove noise other than to raise the detection threshold, our threshold would have to be quite high. Fortunately, we do have an alternative mechanism; we discuss this issue further in Section 5.1.8.

Rather than raising the threshold to 43.8, it seems more prudent to aim for one noise event to exceed threshold in each workunit – a unit of data defined to be 13 s for logistical reasons related to our data processing methods – and sort out false pulses later. (More than one pulse per workunit would be difficult to store in our database.)

In this case, we just want to find C , the number of samples we examine, multiplied by 2 to account for co-adds. This gives us the number of potential pulses (counting all co-adds) per workunit.

$$C = 2^{25} \text{ samples per workunit} \times 14208 \text{ DMs} \times 2 \text{ signs} \times 2 \text{ from co-adds} = e^{28.3}. \tag{42}$$

Since the standard deviation for the exponential is $\sigma = 1$, this means that by setting the threshold at 28.3 we would be looking for pulses that are $28.3 - 1 = 27.3\sigma$ above the mean.

When we co-add $n = 2^\ell$ samples to make one co-added potential pulse, we are adding

up that many exponential distributions. The resulting power has a gamma distribution, with scale parameter 1 and shape parameter n . The pdf is

$$\frac{1}{\Gamma(n)} x^{n-1} e^{-x}. \quad (43)$$

and the complementary cumulative distribution function is defined to be:

$$\int_x^\infty \frac{1}{\Gamma(n)} x^{n-1} e^{-x} dx = \frac{\Gamma(n, x)}{\Gamma(n)}. \quad (44)$$

where $\Gamma(n, x)$ is the upper incomplete gamma function. The first few pdfs are shown in Figure 1.

Then for each $n = 2^\ell$, we want to select a threshold, H_n , such that $\frac{1}{C} = \Gamma(n, H_n)/\Gamma(n)$. We admit pulses of width n only if they have power greater than or equal to H_n . Since a pulse of power H_n would occur by chance with probability $\frac{1}{C}$, and there are precisely C potential pulses in a workunit, we expect to obtain $C \cdot \frac{1}{C} = 1$ false positive per workunit.

Because the probability distribution comes from a sum of identical exponential distributions, we can say that for co-add $n = 2^\ell$, the variance is n times higher, and the standard deviation is \sqrt{n} times higher, than for an exponential. If we define m to be the number of standard deviations of our threshold above the mean, then we are looking for pulses at $m_n = (H_n - n)/\sqrt{n}$. A computation of m_n , using our actual thresholds H_n (as determined by simulation, rather than theory) can be found in Table 1.

3.3.2. Expected discrepancies with the model

A few differences from the model can be expected:

1. **Hydrogen line and filter shape.** We assumed above that the input data is white noise. In practice, this is not the case, because a portion of our band has higher power due to the hyperfine hydrogen line. The strength of this line can vary depending on our RA and dec. Then $\tilde{f}(k)$ no longer have equal standard deviations. This will

cause some correlation between the deconvolved power of adjacent samples, which will modify the pdf of the binned power, increasing the variance.

To see this, consider the simplest, most extreme case: we imagine that the hydrogen line takes the form of a strong delta function in the frequency domain of amplitude A at frequency k_0 , where A is distributed randomly according to a Gaussian distribution with standard deviation σ and mean 0. Then if f_d is the dechirped amplitude in the time domain, $f_d(t) = Ae^{2\pi i k_0 t/N}$. (In other words, the dispersion is not relevant, since the hydrogen line has a single frequency, and we are momentarily assuming it overwhelms the noise.) So $|f_d(t)|^2$ is exponential with power σ^2 , as discussed in Section 3.3.1. Now consider two nearby times $t_1 \sim t_2$ such that the phase of the hydrogen line does not change much between these samples. We want to sum the amplitudes $f_d(t)$ at these nearby times in order to build a co-add. Then $|f_d(t_1) + f_d(t_2)|$ is $A|e^{2\pi i k_0 t_1/N} + e^{2\pi i k_0 t_2/N}|$. Since the phases are similar, this is roughly $2A$, which is Gaussian with standard deviation 2σ and variance $4\sigma^2$. Whereas if we summed samples without the hydrogen line, adding identical and independently distributed (iid) exponentials, the variances would simply add to give $2\sigma^2$. So this model hydrogen line increases the variance.

In actuality, the effect of the hydrogen line is not so pronounced, but the idea is similar. In the same way, the nonuniform shape of our low pass filters also causes the signal to differ from white noise.

2. Other disparities

Even in the absence of the hydrogen line, tests reveal other differences between the theoretical and actual distributions. For high co-adds, the variance is slightly less than expected.

It’s easy to see that power per sample cannot be independently distributed, even in the case of white noise. This is because the total power over all samples must be a constant; in our case, the constant is $32,768 = 2^{15}$, the total number of samples in a FFT. This would certainly result in a smaller variance, but we have not established whether this effect suffices to explain the observed disparity.

3.3.3. Single pulse thresholds: Monte Carlo simulation

To choose our thresholds for the single pulse search, we ran the client on 10 “noise” workunits, which we had constructed to contain only white noise. We kept track of the strongest pulses we found at each co-add. The second largest pulse out of 10 is roughly the 90th percentile, so we set our thresholds at that point. This method gives thresholds that are reasonable as long as we don’t demand that we detect precisely equal numbers of pulses at each co-add. The thresholds suggested by this simulation were within 1% to 5% of the theoretical values.

Because the “noise” workunits contained only white noise, with no Hydrogen line, some of the deviations described above would not be expected to occur. For this reason, our simulation was an imperfect model. However, our concern in this case was not to model the noise and Hydrogen line perfectly, but to obtain a rough estimate for the correct thresholds. The Hydrogen line looks slightly different at different points on the sky, so a perfect simulation would be impossible in any event. As discussed above, our thresholds are supplemented by our RFI and noise rejection algorithms, so we have some flexibility in setting thresholds.

Table 1: Pulse area thresholds H_n , in normalized units such that one sample has an expected power of 1 unit, derived from the Monte Carlo simulation. We denote the implied number of standard deviations above the mean by “ m ”

ℓ	n	H_n	m
0	1	29.1	28.1
1	2	31.6	20.9
2	4	37.9	16.6
3	8	49.4	14.6
4	16	61.3	11.3
5	32	87.0	9.7
6	64	128.9	8.1
7	128	212.6	7.5
8	256	362.0	6.6

3.4. Expected sensitivity

3.4.1. Scattering

We will discuss the sensitivity of Astropulse and other surveys, but we will first consider scattering. Scattering impacts the sensitivity of a survey by limiting its resolution. We suppose that an instantaneous pulse would be broadened by scattering to a width Δt_{sc} . Then we can estimate Δt_{sc} using the empirical formula given in Lorimer & Kramer (2005):

$$\log \Delta t_{\text{sc,ms}} = -6.46 + 0.154(\log \text{DM}) + 1.07(\log \text{DM})^2 - 3.86 \log \nu_{\text{GHz}}. \quad (45)$$

However, this formula applies to sources in the Milky Way. Astropulse spends a substantial fraction of the time looking outside our Galaxy, in which case Δt_{sc} should be much smaller, even for large DMs. But the distribution of the intergalactic medium (IGM)

is not well understood. Lovell et al. (2007) find that extragalactic radio sources at redshifts greater than $z = 2$ do not have microarcsecond structure, suggesting that they are scatter broadened by turbulence in the IGM. This fact can be used to estimate the width Δt_{sc} of the broadened pulse. A microarcsecond of angular broadening at $z = 2$ corresponds to a pulse width of $\Delta t_{\text{sc}} = \theta^2 d/c = 8 \mu\text{s}$ (using $d = 3.57 \text{ Gpc}$). According to Ioka (2003), this is at a DM of about 2000 pc cm^{-3} . So perhaps we can assume that at the (smaller) DMs of our experiment, pulses will have reasonably small widths. For instance, inside the galaxy, a DM of 830 pc cm^{-3} would have a scattering width 400 times smaller than a DM of 2000 pc cm^{-3} . Even if the scattering width were nearly $8 \mu\text{s}$, Astropulse is still good at detecting such pulses. (The threshold is just twice as high as for 1 sample pulses.)

3.4.2. Sensitivity of Astropulse

To calculate Astropulse’s expected sensitivity as a pulse area in $\text{Jy } \mu\text{s}$, we can follow the treatments in Rohlfs & Wilson (2000) and Van Vleck & Middleton (1966), which discuss the effect of “clipping” a noisy analog signal, changing it into a one-bit digital time series.

From these sources we conclude that if F is the flux density and N is the duration in samples of the minimal detectable pulse, then its pulse area is:

$$F \cdot (N t_{\text{sample}}) = \frac{\pi T_0 t_{\text{sample}} (H_N - N)}{2 G}. \quad (46)$$

In this expression:

1. H_N is the power threshold derived from the gamma distribution in Section 3.3.1, and is dependent on N .
2. $T_0 \sim 30 \text{ K}$ is the system temperature.³
3. $G = 10 \text{ K Jy}^{-1}$ is the telescope gain, roughly equal to $\frac{A}{k}$, where:
4. $A = \frac{\lambda^2}{\Omega}$ is the effective area of the telescope and k is Boltzmann’s constant.
5. $\lambda = 21 \text{ cm}$ is the wavelength of the signal.

³http://www.naic.edu/alfa/gen_info/info_obs.shtml

6. $\Omega = 8.1 \cdot 10^{-7}$ is the beam’s solid angle.

So the resulting pulse area is $1.9 (H_N - N) \text{ Jy } \mu\text{s}$.

3.4.3. Sensitivity comparison

While Astropulse detects a signal coherently, other undirected radio surveys use incoherent detection schemes. Typically they use a filter bank, dividing the spectrum into N sub-bands as described in Section 3.2. Deneva et al. (2009) give the sensitivity formula as:

$$S_{\min} = \left(\frac{W}{W_i}\right) \frac{m S_{\text{sys}}}{\sqrt{N_{\text{pol}} B W}}, \quad (47)$$

$$W = (W_i^2 + \Delta t_{\text{DM, ch}}^2 + \Delta t_{\text{DM, err}}^2 + \Delta t_{\text{sc}}^2)^{1/2}. \quad (48)$$

S_{sys} is the system-equivalent flux density.

m is the desired number of standard deviations for the detection threshold,

W is the effective width of the pulse, including broadening due to dispersion and scattering.

W_i is the intrinsic width of the pulse, prior to broadening.

$\Delta t_{\text{DM, ch}}$ is the dispersion within one channel.

$\Delta t_{\text{DM, err}}$ is the error caused by looking at the wrong dispersion measure. We have a DM error of $\frac{1}{2}$ the DM step. (The time error $\Delta t_{\text{DM, err}}$ depends on the bandwidth as well as the DM error.)

$\Delta t_{\text{sc}} \propto f^{-3.86}$ is the error caused by scattering broadening, where f is the pulse’s frequency.

Equation 47 assumes that the channel bandwidth is not so narrow that we are sampling beyond the Nyquist rate. If the channel bandwidth were that narrow, there would be

another contribution to the effective width; but all surveys are careful not to sample beyond the Nyquist rate.

If one cares about the pulse area in Jy μ s rather than the instantaneous flux density, a simpler expression will suffice. If one has observed at a single polarization, the minimum detectable pulse area for that polarization will be $A = \frac{mT_0t}{2G\sqrt{Bt}} = m\frac{T_0}{2G}\sqrt{t/B}$, and the pulse area of both polarizations together will be $A = m\frac{T_0}{G}\sqrt{t/B}$. If both polarizations are observed together and their power is summed, then the minimum detectable pulse area of the two polarizations together is smaller by a factor $\sqrt{2}$, assuming an unpolarized signal.

As discussed in Section 3.4.1, the scattering error is likely to become less important when we are looking outside our Galaxy. So, with an understanding that different conditions would apply when observing the Galaxy, we will ignore the scattering error in Table 2 and set $t = t_{\text{sample}}$. In Table 2, we list characteristics of Astropulse and other surveys and the following conventions are observed:

- m : number of standard deviations for threshold. For any survey that uses incoherent dedispersion, this refers to the standard deviation of a Gaussian distribution, so $m = 6$ suffices to rule out all but 1 in 10^9 spurious detections due to noise. However, for Astropulse, the distribution is a chi-square, as discussed in Section 3.3. In the worst-case scenario – a one-sample pulse – this chi-square is equivalent to an exponential, and $m = 23$ is required to rule out all but 1 in 10^9 spurious detections due to noise. In fact, we selected the value $m = 30$ for Astropulse in the case of one-sample pulses. For all surveys, the optimal value of m is not determined only by the statistical distribution, but also by RFI and noise rejection. For Astropulse, the statistics are discussed at length in Section 3.3.1, and RFI mitigation is discussed in Section 5. In Table 2, we have listed the values of m reported by each survey in available publications. If we could not determine the survey’s m -value, we assumed $m = 6$. We are assuming that astrophysical pulses which surpassed each survey’s stated threshold would be detected by the survey and could be distinguished from noise and RFI.
- t_{sample} : the time resolution of the survey.
- t : the minimum effective duration of a pulse after dedispersion.
- beam Ω : the telescope beam width, in steradians.
- beams: number of simultaneous beams.
- t_{obs} : observation time per beam, in hours.
- N_{pol} : Astropulse detects a pulse using data from a single polarization; many surveys combine two polarizations.

- **sens:** the minimum detectable pulse area. For each survey, the listed sensitivities apply for pulses which are unresolved by that survey. Pulses of duration $0.4 \mu\text{s}$ are unresolved by all surveys, including Astropulse. For pulses of greater duration, Astropulse’s sensitivity will degrade. Table 1, discussed earlier in Section 3.3.1, depicts the relative increase in Astropulse’s minimum detectable pulse area for wider pulses. Other surveys’ sensitivities do not degrade until their time resolution is reached.

In order to compare the surveys using a concrete example, we will also provide information about each survey’s ability to detect evaporating primordial black holes, under specific assumptions: that $M = 10^8$ kg of the black hole’s mass is transformed into a radio signal of bandwidth 1 GHz, and that scattering broadening does not substantially interfere with detection. These assumptions are not intended to be representative of all models of evaporating black holes, and black holes themselves are not the only potential source of microsecond pulses.

- d_{max} : the minimum distance from which an exploding $M = 10^8$ kg black hole would be visible, in kpc. It’s calculated using

$$U_{\text{min}} = \text{energy}/(\text{area} \cdot \text{bandwidth}) = (Mc^2)/(4\pi d_{\text{max}}^2 \cdot 1\text{GHz}), \quad (49)$$

where U_{min} is the pulse area of the minimum detectable signal in Jy μs .

- **rate:** the minimum rate of black hole explosions under which such a black hole would be detectable, $V^{-1}t_{\text{obs}}^{-1}$. Here, $V = (4\pi/3)d_{\text{max}}^3 n_{\text{beams}} \frac{\Omega}{4\pi} = \frac{1}{3}\Omega d_{\text{max}}^3 n_{\text{beams}}$ is the volume of space observed at any one time.

We conclude that Astropulse’s minimum detectable rate (in black holes explosions $\text{pc}^{-3} \text{yr}^{-1}$) is comparable to that of other surveys, but not superior. Astropulse’s rate is similar to Lorimer & Bailes (2007) and Deneva et al. (2009). Our sensitivity to unresolved pulses, in Jy μs , is superior to all other surveys listed except for the Arecibo multibeam survey of Deneva et al. (2009). This sensitivity comes largely from our microsecond time resolution and high gain. Our observation time is also superior. Astropulse does have substantial disadvantages, including a limited bandwidth and narrow ($\Omega = 8.1 \cdot 10^{-7}$) beams.

Table 2: Survey parameters. Parentheses around a value indicate that we assume this value because we could not deduce one from the original paper.

#	author	telescope	year	dedisp	ref	ν_0 (MHz)	m	T_0 (K)	$t_{\text{sample}}(\mu s)$	$t(\mu s)$
1	O’Sullivan et al.	Dwingeloo	1978	incoh	^a	5000	(6)	65	2	2700
2	Phinney & Taylor	Arecibo	1979	incoh	^b	430	6	175	$1.7 \cdot 10^4$	$1.7 \cdot 10^4$
3	Amy et al.	MOST	1989	incoh	^c	843	(6)	-	1	$1.7 \cdot 10^4$
4	Katz & Hewitt	STARE	2003	incoh	^d	611	5	150	125000	125000
5	McLaughlin et al.	Parkes	2006	incoh	^{e, f}	1400	5	21	250	250
6	Lorimer & Bailes	Parkes	2007	incoh	^g	1400	(6)	21	1000	1000
7	Deneva et al.	Arecibo	2008	incoh	^h	1440	5	30	64	64
8	Von Korff et al.	Arecibo	2009	coher	-	1420	30	30	0.4	0.4

#	$\Delta\nu$ (MHz)	G (K Jy ⁻¹)	beam Ω	beams	t_{obs} (h)	N_{pol}	sens (Jy μs)	d_{max} (kpc)	rate (pc ⁻³ yr ⁻¹)
1	100	0.1	$6.6 \cdot 10^{-6}$	1	46	1	$2 \cdot 10^4$	61	$3.8 \cdot 10^{-7}$
2	16	27	$6.6 \cdot 10^{-6}$	1	292	1	1300	240	$9.4 \cdot 10^{-10}$
3	3	-	$3.6 \cdot 10^{-8}$	32	4000	1	$1.6 \cdot 10^5$ ⁱ	22	$5.6 \cdot 10^{-7}$
4	4	$6.1 \cdot 10^{-5}$	1.4	1	13000	2	$1.5 \cdot 10^9$	0.22	$1.3 \cdot 10^{-7}$
5	288	0.7	$1.3 \cdot 10^{-5}$	13	1600	2	99	870	$1.5 \cdot 10^{-13}$
6	288	0.7	$1.3 \cdot 10^{-5}$	13	480	2	240	560	$1.8 \cdot 10^{-12}$
7	100	10	$8.1 \cdot 10^{-7}$	7	420	2	8.5	3000	$4.2 \cdot 10^{-13}$
8	2.5	10	$8.1 \cdot 10^{-7}$	7	1540	1	55	1200	$1.9 \cdot 10^{-12}$

^aO’Sullivan et al. (1978)

^bPhinney & Taylor (1979)

^cAmy et al. (1989)

^dKatz et al. (2003)

^eMcLaughlin et al. (2006)

^fManchester et al. (2001)

^gLorimer & Bailes (2007)

^hDeneva et al. (2009)

ⁱMOST has 1 mJy of noise in each beam after 12 hours, <http://www.physics.usyd.edu.au/sifa/Main/MOST>

4. Distributed computing : the BOINC platform

Astropulse runs on the BOINC platform (Anderson 2004), an acronym for “Berkeley Open Infrastructure for Network Computing.” BOINC is a set of programs that organizes volunteers’ home computers to perform scientific calculations. In a typical BOINC project, a researcher has a computing problem that can run in parallel, that is, on several machines at once. Perhaps the problem involves searching a physical space (for Astropulse, this space is the sky), and performing the same computation on each point in that space (for Astropulse, this computation is dedispersion.) The first BOINC project, SETI@home, searched the sky for narrowband transmissions. The space could also be a parameter space, for instance a space of potential climate models (climateprediction.net) or protein configurations (Rosetta@home). The visible manifestation of a BOINC project is an informative screen saver. Figure 2 shows the Astropulse screen saver.

Although the volunteers are providing their computers for free, the bandwidth and storage space required to distribute data to the volunteers is not free. A project is suitable for BOINC only if it is computation-intensive. That is, the monetary cost to perform the computation must be greater than the monetary cost of distributing the data. Coherent dedispersion satisfies this requirement, because we must perform FFTs at many DMs.

The researcher for a BOINC project need not be affiliated with UC Berkeley, or with the BOINC development team at Berkeley (although we happen to be so affiliated). BOINC is open source, and can be downloaded, compiled, and operated by anyone with sufficient technical skills; about 50 projects currently exist outside Berkeley.

Likewise, volunteers need not have any particular technical knowledge. They just have to navigate to the BOINC web page with their web browser, and follow the instructions to download the Astropulse “client” program. Astropulse has access to around 500,000 volunteers, each of whose machines might have 2 GFLOPs of processing power, and be on 1/3 of the time, for a total of 300 TFLOPs – as much as the world’s fastest general purpose supercomputer in 2007, IBM’s Blue Gene / L⁵. Since that time, the processing power of the fastest supercomputer has increased to 8000 TFLOPs or more⁶.

⁵<http://www.top500.org/list/2007/06/100>

⁶<http://www.top500.org/list/2011/06/100>

5. RFI mitigation

The current section describes the methods by which we have classified all detected pulses, deciding whether they might correspond to RFI or noise. We will begin by discussing a number of RFI and noise mitigation methods. We will then define a figure of merit statistic for each method, and use Monte Carlo simulations to argue that astrophysical signals will not be excluded by our RFI and noise mitigation methods. Finally, we will mention our successful detection of giant pulses from the Crab pulsar, providing further evidence that astrophysical signals will not be excluded.

The dominant RFI sources at Arecibo Observatory are nearby radars, which emit one of (at least) 6 repeating patterns. These include the Federal Aviation Administration (FAA) radar, used for air traffic control around Puerto Rico; the aerostat radar, used for drug interdiction; and others.

5.1. RFI and noise mitigation methods

We rejected RFI and noise using several methods. We will first discuss methods implemented prior to pulse detection. Most of these methods (5.1.2, 5.1.3, and 5.1.4) blanked segments of the data which were likely to be contaminated by radio frequency interference. We will then describe a number of post-detection techniques. Using these techniques, we characterized pulses as probable RFI by examining their dispersion measures, polarizations, shapes, and other properties, and by examining other pulses detected at nearby times.

5.1.1. *Arecibo’s high pass filter*

Arecibo can turn on a high pass filter in the receiver that will reject the FAA radar’s band from the data. However, Astropulse operates commensally, and our partner surveys usually require that this filter is turned off.

5.1.2. *Hardware blanker*

Arecibo Observatory provides us with a blanking signal (which we will call the “hardware” blanking signal), a single bit which is turned on when the FAA radar is transmitting a pulse. The “hardware” blanker has two components:

1. The hardware component at Arecibo, which adds the blanking bit to our tape files.
2. A software component, activated at a later point in our data processing pipeline, which detects the bit and blanks the appropriate data. (This software component still counts as a part of our “hardware” blanker.)

It is critical that we blank the data using noise that has the same frequency profile as the clean data. If we instead blank the data using white noise (bits set randomly to 1 and 0) artificial signals will be introduced, since the white noise does not match the rest of the data. Therefore, we need to blank the data by replacing it with noise whose frequency envelope matches that of our data recorder’s 2.5 MHz bandpass filter.

Unfortunately, the hardware blanker is imperfect. First, we believe it doesn’t mark every FAA radar pulse. Sometimes the radar’s phase changes, and it takes some time for the hardware blanker to catch up. At other times, a single radar pulse may arrive that is out of sync with the other pulses. Second, the hardware blanker only searches for the FAA radar, not for other radar. So we have written our own software blanker, which processes the data downstream from the hardware blanker. The software blanker handles both the FAA radar and the aerostat radar.

5.1.3. *Software blanker*

The software blanker is an algorithm that runs on a computer in Space Sciences Lab. It examines the data for the repeating patterns that signify radar. It looks specifically for the FAA and aerostat radar.

To find either radar, the blanker looks at samples in groups of 10. A radar pulse would consist of samples where the bits are predominantly 1 or predominantly 0. (That’s 280 bits, counting all 28 bits for each sample, 2 polarizations, 7 beams, and both real and imaginary bits.) At maximum strength, the radar will produce long strings of bits that are all set to 1, regardless of whether they represent real or imaginary data. At lesser strengths, the radar

produces less skewed sets of samples, with a “ring down” oscillation between 1 and 0 bits. Nevertheless, these sets of samples at lesser strengths can provide an important indication of radar.

The blanker folds the data over 25 seconds at the known radar periods, which are 35,262 samples for the FAA radar, and 57,571 samples for the aerostat radar. (For each radar source, the intervals between pulses vary according to a set sequence. So the stated FAA period actually contains 5 pulses, and the aerostat period contains 7 pulses.) Actually, we fold at around 200 trial periods, each varying slightly from the average radar period. This is necessary because the radar’s period can drift slightly. We threshold the resulting amplitudes at 25% above the mean.

If a radar signal has been detected, we blank data at regular intervals over the 25 seconds, accounting for the varying interpulse periods. Typically, the radar shape is square: a total radar silence, followed by radar of duration 800 samples, followed by more silence. However, we have observed that individual radar pulses may be smeared by up to 50 samples on either side of the region of duration 800. Therefore, we blank 100 samples on either side, for safety. In total, we typically blank 1,000 samples for each radar pulse.

Most of our data contains FAA radar signals, and relatively little has aerostat signals. Overall, we blank about 14% of our data during this stage. The software blanker is very effective in comparison with the hardware blanker; more than 99% of pulses that are detected by the hardware blanker are also detected by the software blanker, whereas the reverse is not true.

5.1.4. *Client blanker*

The software blanker is effective at removing FAA and aerostat radar patterns from our data, but other types of RFI exist which do not have these periods. In order to detect and remove these other signals, we implement a blanker in the Astropulse client. Like the software blanker, the client blanker searches for RFI that is strong enough to saturate our electronics, producing a long string of identical samples. We have found that this string of identical samples not only indicates RFI at that particular instant, but also warns that other nearby data in the same workunit may be contaminated by RFI. Although a string of identical samples is not itself a dispersed pulse, it can signify that dispersed RFI may be present in the nearby data. Astropulse could record this dispersed RFI as candidate pulse in our database. Therefore, we blank all data within 400,000 samples (0.16 s) of the detected event. This figure, 400,000 samples, was determined empirically, by examining many workunits showing strings of identical samples. We found that these strings were often accompanied by dispersed pulses within a distance of 400,000 samples. We do not

believe that these dispersed pulses represented true astrophysical transients, because they were very common, and because, at least for our system, the strings of identical samples are a known feature of RFI.

The client blanker differs from the software blanker in that we consider individual RFI events, rather than folding several events together. This enables detection of RFI with unknown periods. The client blanker proceeds by performing a Fourier transform of a segment of the data, and examining the power in the central bin (the DC component.)

5.1.5. Fraction blanked restriction

This and all following methods occurred in the post-detection phase; that is, we stored pulses in our database prior to determining whether they passed or failed this test. For the current method, we consider each workunit and record the fraction of the data that we blanked using the client blanker. We remove workunits entirely if too much RFI was present, since we have observed that the presence of too much RFI in one region of the workunit may indicate some amount of RFI in other regions.

5.1.6. DM repetition

If we see a signal at the same DM repeatedly over a short period of time and at different parts of the sky, we conclude that it came from a terrestrial source and reject these signals as RFI. There is no reason that the same DM should have been observed from several different directions in quick succession unless the source was terrestrial. This test, like the multi-beam and multi-polarization tests described below, is implemented by a program that examines our database of detected candidate pulses, putting them in time order and searching for the appropriate pattern.

5.1.7. Multiple simultaneous beams

This method, which makes use of the ALFA receiver’s multiple beams, proved to be less useful than we had hoped. We did not make use of it for the calculations discussed in this paper, but nevertheless we describe the method here. The method relies on the fact that our beams are separated by several arcminutes, so an astrophysical source (or any

relatively weak source) should not appear in multiple beams simultaneously. However, a very strong source, for instance a terrestrial source, might appear in the beams’ sidelobes. The source’s radio waves might arrive at the telescope by scattering from nearby terrain, or by bouncing off the telescope support structure. In this case, the waves might appear in multiple beams. Since the telescope never points at or below the horizon, such sources would appear only in the sidelobes and never in the main lobe.

Therefore, we could rule out some RFI by ignoring pulses that appear in multiple beams simultaneously. Unfortunately, experiments show that a few real, astrophysical signals appear in multiple beams simultaneously – namely, the detected Crab pulses. This may happen because the main lobes intersect slightly, albeit at greatly diminished sensitivity, or because strong astrophysical signals could be detected in sidelobes. So such a method is imperfect at best.

5.1.8. Two simultaneous polarizations

A signal from an unpolarized astrophysical source will appear in both polarizations simultaneously (unless it is only marginally detectable.) RFI might also behave this way, but noise will not. Therefore, we can reject a great deal of noise by requiring detections in two simultaneous polarizations. Unfortunately, highly polarized astrophysical signals may also be rejected, especially if the signal’s axis of polarization lines up with the telescope’s axis of polarization. This drawback is balanced by the extraordinary efficacy of the polarization test as a noise rejection technique.

5.1.9. Frequency profile

We are looking for broadband pulses with a short intrinsic timescale. Thus, the pulse should have roughly the same mean power at all frequencies. We perform a chi-square test to determine whether the mean power is the same everywhere. However, the chi-square distribution is not a perfect description of the power vs. frequency distribution unless the power has a Gaussian distribution at each frequency. In fact, the power should have an exponential distribution, not Gaussian.

The frequency profile test calculates a “log_prob” statistic, which is the natural log of the estimated probability that this frequency profile would occur by chance. However, if the chi square is inaccurate, so is the log_prob. Nevertheless, the log_prob should decrease (and is negative) as the power becomes concentrated at particular frequencies. Although

the `log_prob` has uncertain meaning in the absolute sense, its relative value is meaningful.

5.2. Figure of merit

We can assign a figure of merit to each RFI rejection algorithm, or to all algorithms together. The figure of merit is defined as:

$$(\% \text{ of astrophysical pulses passing}) / (\% \text{ of all candidate pulses passing}).$$

A candidate pulse is a detection that is above threshold, which may or may not correspond to a signal from an astrophysical source. The purpose of an RFI rejection algorithm is to throw out spurious candidate pulses, while hopefully preserving most of the astrophysical pulses. If the figure of merit is equal to 1, the algorithm does not change the percent of candidate pulses that are astrophysical. In other words, we could have achieved the same result by throwing out a random collection of our pulses. Therefore, an algorithm cannot be useful unless its figure of merit is greater than 1.

This definition of the figure of merit is not the only one imaginable. For example, suppose we have 1,000 pulses, of which 100 are astrophysical, and two algorithms. The first algorithm cuts the list down to 10 pulses, of which 9 are astrophysical. It has a figure of merit equal to 9. The second algorithm instead cuts the list down to 100 pulses, of which 80 are astrophysical. It has a figure of merit equal to 8. Then one might prefer the latter algorithm, on the grounds that it yields more data to work with (even though a smaller percent of that data is good).

Nevertheless, this figure of merit is reasonable, and we calculate its value for each of our RFI rejection algorithms in the following sections. In our calculations, we will assume that the vast majority of candidate pulses are due to noise, since if this large number of transient radio signals were of astrophysical origin, the phenomenon would have been reported by other surveys.

In the following sections, we will compute a figure of merit for the post-processing methods. We cannot compute a figure of merit for pre-processing methods, since we cannot count the number of candidate pulses (such as RFI) that we would have detected in absence of those methods.

5.2.1. *Fraction blanked restriction: figure of merit*

Empirically, it turns out that we can obtain the best figure of merit by passing only those workunits for which the fraction blanked (by the client blanker) is $< 20\%$. Note that the client blanker has already removed a portion of each workunit. Here, we do not consider the figure of merit resulting from the operation of the client blanker itself. Rather, we are throwing out workunits for which a large portion has already been removed. As of December 2009, the figure of merit statistics are as given in Table 3. Instead of simulating astrophysical pulses, we have counted the space available for such pulses in all workunits. Only unblanked space may contain astrophysical pulses (or any pulses that originate outside the telescope), and we are assuming that the likelihood for an astrophysical pulse to appear in a workunit is proportional to the amount of unblanked space in that workunit.

5.2.2. *DM repetition: figure of merit*

To simulate the fraction of astrophysical pulses that would be accepted by the DM repetition algorithm, we performed a Monte Carlo study, generating a list of 37,572 pulses at random times. We made use of this list of pulses to analyze the DM repetition, multiple beams, and simultaneous polarizations tests. The random times were determined by considering the start times of actual workunits, then selecting a random time within that workunit. Other relevant parameters, such as dispersion measure, beam, and polarization, were assigned a uniformly distributed random value within the allowed range. (Pulse area was not relevant for this list of test pulses.)

Using the random dispersion measures, we counted the number of detected pulses with the same dispersion measure preceding and following the test pulses. The test pulses were accepted or rejected using the same criteria as the DM repetition RFI rejection method.

Monte Carlo statistics for simulated astrophysical pulses, after 3 passes through 12,524 workunits, generating 37,572 test pulses, are listed in Table 3.

5.2.3. *Multiple simultaneous beams: figure of merit*

To simulate the fraction of astrophysical pulses that would be accepted by the “simultaneous beams” algorithm, we used the same Monte Carlo study that we performed for DM repetition. The test pulses were accepted or rejected using the criteria from the

“simultaneous beams” RFI rejection method.

Monte Carlo statistics, for the same pulses as in Section 5.2.2, are listed in Table 3.

5.2.4. Two simultaneous polarizations: figure of merit

If all detected astrophysical pulses were completely unpolarized, or were above threshold in both polarizations, then all of them would pass the “simultaneous polarizations” test. However, even if the astrophysical component of the pulse is unpolarized, the noise component is independent of the astrophysical component. Thus, pulses near threshold may be detectable in only one polarization.

To simulate the fraction of unpolarized astrophysical pulses that would be accepted by this test, we generated pulse area values with a cumulative distribution $c(s) \propto s^{-3/2}$, or a probability density function $h(s) \propto s^{-5/2}$, where s , drawn from the random variable S , is the pulse area.

The reason for the $s^{-3/2}$ cumulative distribution is that if we assume a standard candle source (same luminosity vs. time for all sources), then the sources at distance r have flux at Earth proportional to $\frac{1}{r^2}$. The number of sources within distance r (hence with flux greater than $S \propto \frac{1}{r^2}$), is proportional to $r^3 \propto S^{-3/2}$.

After determining the test pulse’s area, we generate two mini workunit files that contain the pulse. Each file combines the pulse with noise randomly, so that different noise is generated in the two mini workunits. Then, we dedisperse the two files and find the noise-modified pulse areas. The pulse passes the “simultaneous polarization” test if it is above the detection threshold in both polarizations. If it is only above threshold in one polarization, it fails the test. And if it is below threshold in both polarizations, it would not be detected at all, so it does not pass or fail.

Note that for a given pulse area threshold, there is a unique probability density function (pdf) with $h(s) \propto s^{-5/2}$, so there is no ambiguity about normalization. If more astrophysical sources are present, the total number of sources detected will increase, but the pdf will not change.

After 1,000 pairs of test pulses, the figure of merit statistics are given in Table 3. The x statistic is not the number of pulses generated (2,000), but the number detected; some pulses were below threshold. In the table, the x statistic counts pairs of corresponding

pulses as two, whereas the y statistic counts each pair as a single pulse, and z is their ratio. The z_2 statistic was arrived at in a similar manner.

5.2.5. Frequency profile: figure of merit

Using the same mini workunits generated for the polarization test, we determine whether the pulse would pass the frequency profile test. The pulse passes the frequency profile test if its spectrum is flat. Again, the pdf of the pulse area is unique, given the pulse area threshold, therefore there is no ambiguity as to the pulse powers we should use.

After a Monte Carlo using threshold $\log_prob > -1$, and after 1,000 pairs of test pulses, the figure of merit statistics are given in Table 3. The x statistic is not the number of pulses generated (2,000), but the number detected; some pulses were below threshold.

5.2.6. Overall: figure of merit

There seems to be no reason that the astrophysical pulses' passing fractions, as described above, should be correlated. (Especially if we exclude the multi-beams test, which is probably unreliable.) An astrophysical pulse that passes the DM repetition test is no likelier than any other to pass the multi-pols test, the fraction blanked test, or the frequency profile test.

To see this, one has to consider the tests in pairs, and think about the nature of the tests. In each case, the property measured by one test is entirely unrelated to the property measured by the other. A pulse passes the multi-pols test if it is strong and/or unpolarized, and it fails the DM repetition test if nearby (noise or RFI) pulses have the same DM as the signal. It passes the fraction blanked test if its workunit has a lot of RFI that overwhelms the receiver or IF electronics, and it passes the frequency profile test if its spectrum is flat. So we will assume that the four tests are statistically independent, an assumption which we discuss in the following section.

So we expect the fraction of astrophysical pulses passing all tests to be: $0.557 \cdot 0.958 \cdot 0.481 \cdot 0.797 = 0.205$, where we have just multiplied the fraction passing from each test above.

On the other hand, the fraction of candidate pulses passing all tests is: $47/412001 =$

0.000114, as of February 2010. (Not counting pulses from our observation of the Crab pulsar.) This makes for a figure of merit equal to 1797, substantially larger than the product of the individual figures of merit, which is 33 (see Table 3.) This makes sense, because the multi-pols test is designed to catch noise, whereas the other tests are designed to catch RFI. So we might expect each test to be less effective on its own, but more effective in combination with other tests. (For instance, imagine a fictitious data set in which 49% of all signals are noise, 49% are RFI, and 2% are real. If algorithm A removes all noise, and algorithm B removes all RFI, then the two together have a figure of merit of $1/0.02 = 50$, whereas separately they have $1/0.51 \approx 2$.)

5.2.7. *Statistical independence of our tests*

In Section 5.1, we have considered several RFI and noise rejection methods. Our four post-detection methods were as follows:

- A. Fraction blanked restriction
- B. DM repetition
- C. Two simultaneous polarizations
- D. Frequency profile

We also described how we have tested the figure of merit for each of these four methods individually, using Monte Carlo simulations. In each case, we attempted to make reasonable assumptions about the distribution of astrophysical and spurious signals that would be received; these assumptions are detailed in the relevant sections above. In the case of methods (C) and (D), we were required to model the power of the test pulse relative to the noise. For methods (A) and (B) it was not necessary to model the power, since these methods depended only on the arrival time and dispersion measure of the pulse. For each of the four RFI and noise rejection methods, we have computed, using a simulation, the fraction of astrophysical pulses that would pass the test; see Table 3.

However, ultimately we are not interested in the individual algorithms, but in the overall effect of our entire pipeline. One might be concerned that astrophysical pulses will pass individual algorithms, but will fail when they are presented with multiple algorithms. Given the four individual passing fractions just mentioned, can we derive the overall passing fraction for astrophysical pulses?

To answer this question, we must determine whether our RFI rejection methods are statistically independent. That is, if method (A) allows astrophysical signals to pass with probability p_A , and method (B) allows signals to pass with probability p_B , then the two methods are statistically independent if the probabilities for the joint outcomes are as given in Table 4.

Another way of saying this is that two methods are independent if their outcomes are uncorrelated. In the event that all four methods are jointly independent, then the joint passing fraction can be found by multiplying the individual passing fractions.

However, if all methods are not independent, the joint passing fraction may not be the same as the product of the individual passing fractions. In particular, we would like to make sure that the joint passing fraction for astrophysical pulses is not too low. Therefore, we should attempt to verify that the four methods are statistically independent. We might attempt to accomplish this by simulating the four methods using computer generated fake “astrophysical” signals. Unfortunately, it turns out that the statistical independence of the four methods is highly contingent on our choice of a model for the behavior of the astrophysical pulses and of the RFI. We do not observe the astrophysical pulses directly, nor do we measure all properties of the RFI. Therefore, we must make some assumptions about their properties, and different assumptions will lead to different models of the four RFI rejection methods.

There is no way out of this: with incomplete information about the astrophysical pulses and RFI, we must make some assumptions. We would not be able to run a simulation to test these assumptions, since any such simulation would depend upon the very assumptions it was supposed to test. Therefore, the best we can do is to carefully state the assumptions that lead us to model the four methods as being statistically independent. To do this, we consider the four methods in pairs, and explain what it would mean to assume that each pair is statistically independent.

When considering the pairs, we are only concerned about factors that would produce negative correlations, where pulses which pass one test are likely to fail another. These negative correlations would reduce the number of astrophysical pulses that would make it through our algorithms. (Positive correlations, where pulses which pass one test are likely to pass another, would actually improve our ability to detect astrophysical pulses.)

As we consider pairs of methods, our language will be somewhat repetitive, since each method will be discussed three times. The reader may wish to examine just one or two entries, to get a flavor of the kind of reasoning that is required.

(A) and (B): fraction blanked and DM repetition:

Astrophysical pulses will fail method (A) if the client blanker detects a large amount of RFI at a nearby time. They will fail method (B) if nearby pulses have the same dispersion measure as the astrophysical pulse. This would be most likely to happen if the RFI over a particular period of time was concentrated at a particular dispersion measure. Thus, the two methods will be negatively correlated if repeating dispersion measures in the RFI are correlated with low amounts of RFI, and nonrepeating dispersion measures are correlated with high amounts of RFI. We are unaware of any such effect.

(A) and (C): fraction blanked and simultaneous polarizations

Astrophysical pulses will fail method (A) if the client blanker detects a large amount of RFI at a nearby time. They will fail method (C) if they are too weak to show up in both polarizations, or if they are highly polarized. The two methods will be negatively correlated if weak or polarized astrophysical pulses are more likely to arrive during periods of few RFI detections. Since astrophysical pulses come from a different source than RFI, this seems unlikely.

(A) and (D): fraction blanked and frequency profile

Astrophysical pulses will fail method (A) if the client blanker detects a large amount of RFI at a nearby time. They will fail method (D) if they are not broadband, or (more likely) if their profile is distorted by noise or RFI to appear as if it is not broadband. The two methods will be negatively correlated if narrowband RFI, of a sort that could distort an astrophysical signal's frequency profile, is correlated with low amounts of RFI. We are unaware of any such effect.

(B) and (C): DM repetition and simultaneous polarizations

Astrophysical pulses will fail method (B) if nearby pulses have the same dispersion measure as the astrophysical pulse. They will fail method (C) if they are too weak to show up in both polarizations, or if they are highly polarized. The two methods will be negatively correlated if weak or polarized astrophysical pulses are less likely to arrive during periods of RFI with repeating dispersion measure. Since astrophysical pulses come from a different source than RFI, this seems unlikely.

(B) and (D): DM repetition and frequency profile

Astrophysical pulses will fail method (B) if nearby pulses have the same dispersion measure as the astrophysical pulse. They will fail method (D) if they are not broadband, or (more likely) if their profile is distorted by noise or RFI to appear as if it is not

broadband. The two methods will be negatively correlated if RFI with repeating dispersion measure is correlated with low amounts of narrowband RFI, of a sort that could distort an astrophysical signal’s frequency profile. We are unaware of any such effect.

(C) and (D): simultaneous polarizations and frequency profile

Astrophysical pulses will fail method (C) if they are too weak to show up in both polarizations, or if they are highly polarized. They will fail method (D) if they are not broadband, or (more likely) if their profile is distorted by noise or RFI to appear as if it is not broadband. The two methods will be negatively correlated if weak or polarized astrophysical pulses are less likely to arrive at the same time as narrowband RFI. Since astrophysical pulses come from a different source than RFI, this seems unlikely.

Note that these six pairs do not exhaust all of the theoretical possibilities; it is possible for three quantities to be statistically dependent even if two are independent; however, similar sorts of reasoning apply to the case of three quantities’ simultaneous correlations.

5.2.8. *Detection of giant pulses from the Crab pulsar*

The reasoning in Section 5.2.7 aimed to demonstrate that astrophysical pulses have a good chance to get through our RFI mitigation algorithms. Additional supporting evidence comes from our detection of several giant pulses from the Crab.

We observed the Crab pulsar intermittently from 11:39 AST to 13:27 AST on June 7, 2009. We detected 1,404 candidate pulses over that time period, of which 89 passed all of our RFI and noise mitigation tests. All 89 pulses had dispersion measures ranging from 50.0 pc cm^{-3} to 62.6 pc cm^{-3} , with a standard deviation of 1.8 pc cm^{-3} . This tells us that Astropulse can typically estimate a dispersion measure to within 1.8 pc cm^{-3} , or within 6 pc cm^{-3} in the worst case. Most of these pulses seem to be duplicate representations of a single astrophysical pulse (or group of pulses). Astropulse may detect the same pulse in different beams, polarizations, or scales. The 89 pulses passing our tests corresponded to 3 groups clustered in time, indicating that we detected at least 3 giant pulses.

Of the 1,404 candidate pulses, 171 had dispersion measures ranging from 50.0 pc cm^{-3} to 62.6 pc cm^{-3} . These candidate pulses clustered into 10 groups. We can use this fact to estimate the passing fraction of astrophysical pulses from the Crab. At most 10 pulses were astrophysical, therefore the passing fraction was at least $3/10 = 30\%$, which corresponds roughly to our simulated value of 0.205.

We detected these same pulses independently, using a spectrometer at Arecibo Observatory, and verified that the two methods found the pulses at precisely the same times. Therefore, we can be confident that we have detected giant pulses from the Crab pulsar, indicating that Astropulse can detect at least some astrophysical signals.

6. Conclusion

We designed software to search for microsecond transient radio pulses, using coherent dedispersion to examine the microsecond timescale. We have focused on evaporating primordial black holes as a potential source of microsecond pulses, but other sources are possible, including giant pulses from pulsars, RRATs, and as yet unknown astrophysical phenomena. In order to obtain the computational power required for coherent dedispersion, we distributed the software to around 500,000 volunteers, and processed 1,540 hours of observation time at the Arecibo telescope using the Arecibo L-band Feed Array (ALFA) receiver. We observed simultaneously with each of 7 beams and 2 polarizations per beam, for a total of 21,600 hours of data.

In this paper, we have presented the design of the Astropulse experiment, including the scientific rationale and goals of the project, details of the algorithm, data processing, and distributed computing methodology, and our techniques for RFI and noise mitigation. We compared Astropulse’s sensitivity to other searches by computing the minimum detectable pulse area, in $\text{Jy } \mu\text{s}$, for a microsecond pulse that could be detected by each survey. In this respect, Astropulse is more sensitive than all other radio surveys we considered, except for that of Deneva et al. (2009). Astropulse is able to detect pulses of area $54 \text{ Jy } \mu\text{s}$, whereas Deneva et al. (2009) can detect pulses of area $8.5 \text{ Jy } \mu\text{s}$. Astropulse also had the second-largest amount observation time, a total of 10,800 hours of data counting one polarization, whereas McLaughlin et al. (2006) observed for 20,800 hours.

We employed multiple techniques to mitigate radio frequency interference (RFI). Several layers of hardware and software were devoted to blanking known terrestrial radio sources in the vicinity of the Arecibo telescope. We ruled out signals which appeared at the same DM repeatedly at multiple points on the sky, since an astrophysical source would be located at a single point. We also ruled out signals which did not occur in both polarizations simultaneously, searching only for unpolarized signals. Although we lost the opportunity to detect polarized signals, we found that this polarization criterion allowed us to reject a substantial amount of noise. Finally, we required signals to be broadband, having a similar power at all frequencies, a condition that would hold true for any signal with a short intrinsic timescale. By combining these criteria, we were able to reject all but approximately 1 in 10,000 detected pulses. By performing simulations, we estimated that about 1 in 5 astrophysical pulses would pass our RFI rejection criteria. Therefore, as long as we have detected 5 or more astrophysical pulses, we would expect some to remain after

the RFI mitigation step.

In a future paper, we will describe the results of the Astropulse experiment, outlining the characteristics of the pulses that remained, and our conclusions about them.

Table 3: Figures of merit for RFI mitigation algorithms. In the table, x is the number of pulses analyzed in a Monte Carlo test or other simulation, y is the number of those pulses that pass this test, and z is the fraction passing for simulated pulses. x_2 is the number of candidate pulses analyzed so far, y_2 is the number of candidate pulses passing, and z_2 is the fraction passing for candidate pulses. The figure of merit (FoM) is defined as z/z_2 . Note that in the row for “fraction blanked,” x and y refer to un-blanked space in all workunits (before and after the test is applied), in units of full workunit lengths.

algorithm	x	y	z	x_2	y_2	z_2	FoM
fraction blanked	2,457,187	1,368,632	0.557	256,085	122,627	0.479	1.16
DM repetition	37,572	35,994	0.958	204,994	114,795	0.560	1.71
multi-beams	37,572	37,420	0.996	256,085	220,573	0.861	1.16
multi-pols	1,086	522	0.481			0.0386	12.5
frequency profile	1,075	857	0.797	246,870	149,277	0.605	1.32

Table 4: Criterion for the statistical independence of two RFI mitigation methods. Two methods are independent if their outcomes on astrophysical pulses are statistically uncorrelated, so that the probability of each outcome pair is given by the expression in the table.

Outcome	Probability
pass A and pass B	$p_A \cdot p_B$
pass A and fail B	$p_A \cdot (1 - p_B)$
fail A and pass B	$(1 - p_A) \cdot p_B$
fail A and fail B	$(1 - p_A) \cdot (1 - p_B)$

We are grateful to the National Science Foundation (Grant AST-0808175), NASA (Grant NNX09AN69G), and the Friends of SETI@home for funding Astropulse and the SETI@home servers. The Arecibo Observatory is the principal facility of the National Astronomy and Ionosphere Center, which is operated by the Cornell University under a cooperative agreement with the National Science Foundation. We would like to thank Luke Kelley for his work on the software radar blanker. JV would like to thank his wife, Lucy, for her support and encouragement. PBD is a Jansky Fellow of the National Radio Astronomy Observatory.

REFERENCES

- Amy, S. W., Large, M. I., & Vaughan, A. E. 1989, *Proc. Astron. Soc. Aust.*, 8, 2
- Anderson, D. P. 2004, in *Proceedings of the 5th IEEE/ACM International Workshop on Grid Computing*
- Carter, B., Gibbons, G. W., Lin, D. N. C., & Perry, M. J. 1976, *A&A*, 52, 427
- Cordes, J. M. 2008, in *AIP conference proceedings*, Vol. 983, 40 years of pulsars: millisecond pulsars, magnetars and more, 567
- Cordes, J. M., & Lazio, T. J. W. 2003, *arXiv:astro-ph/0207156v3*
- Deneva, J. S., Cordes, J. M., et al. 2009, *ApJ*, 703, 2259
- Guélin, M. 1973, *Proc. IEEE*, 61, 1298
- Hagedorn, R. 1965, *Nuovo Cimento Suppl.*, 3, 147
- Hankins, T. H., Kern, J. S., Weatherall, J. C., & Eilek, J. A. 2003, *Nature*, 422, 142
- Hankins, T. H., Stineberg, D. R., & Rawley, L. A. 1987, *Astrophysical Journal*, 315, 149
- Hawking, S. 1971, *MNRAS*, 152, 75
- . 1974, *Nature*, 248, 30
- Ioka, K. 2003, *ApJ*, 598, L79
- Ivezic, Z., Tyson, J. A., Allsman, R., Andrew, J., & Angel, R. 2008, <http://arxiv.org/abs/0805.2366>
- Kaiser, N. 2004, *Proc. SPIE*, 5489, 11
- Katz, C. A., Hewitt, J. N., Corey, B. E., & Moore, C. B. 2003, *PASP*, 115, 675
- Kowalski, M., & Mohr, A. 2007, *Astropart. Phys.*, 27, 533
- Lattimer, J. M., Prakash, M., Masak, D., & Yahil, A. 1990, *ApJ*, 355, 241
- Levinson, A., Ofek, E. O., Waxman, E., & Gal-Yam, A. 2002, *ApJ*, 576, 923
- Lorimer, D., & Bailes, M. 2007, *Sci*, 318, 777
- Lorimer, D., & Kramer, M. 2005, *Handbook of Pulsar Astronomy* (University Press, Cambridge)
- Lovell, J. E. J., et al. 2007, in preparation, <http://arxiv.org/abs/astro-ph/0701601>
- MacGibbon, J. H., Bailes, M., & Weber, B. R. 1990, *Phys. Rev. D*, 41, 3052

- Manchester, R. N., et al. 2001, MNRAS, 328, 17
- McLaughlin, M. A., Lyne, A. G., Lorimer, D. R., Kramer, M., Faulkner, A. J., Manchester, R. N., Cordes, J. M., & Camilo, F. 2006, Nature, 439, 817
- O’Sullivan, J. D., Ekers, R. D., & Shaver, P. A. 1978, Nature, 276, 590
- Peek, J. E. G., & Heiles, C. 2008, arXiv:0810.1283
- Phinney, S., & Taylor, J. H. 1979, Nature, 277, 117
- Popov, M. V., & Stappers, B. 2007, A&A, 470, 1003
- Raine, D., & Thomas, E. 2005, Black Holes: An Introduction (Imperial College Press, London)
- Rees, M. J. 1977, Nature, 266, 333
- Rohlfs, K., & Wilson, T. L. 2000, Tools of Radio Astronomy (Springer-Verlag, Berlin)
- Sallmen, S., Backer, D. C., Hankins, T. H., Moffett, D., & Lundgren, S. 1999, ApJ, 517, 460
- Stanimirovic, S., & Putnam, M. 2006, ApJ, 653, 1210
- Ukwatta, T. N., et al. 2010, in preparation, <http://arxiv.org/abs/1003.4515>
- Van Vleck, J., & Middleton, D. 1966, Proc. IEEE, 54, 2
- Vestrand, W. T., et al. 2005, Nature, 435, 178




Fig. 1.— Gamma distributions. The x axis is integrated power (divided by 2, as per our convention), and the y axis is probability per unit power. The leftmost distribution, which is exponential, belongs to $n = 1$. The rest are $n = 2, 4, 8, 16, 32, 64$. Notice that the rightmost distribution is nearly a normal distribution.

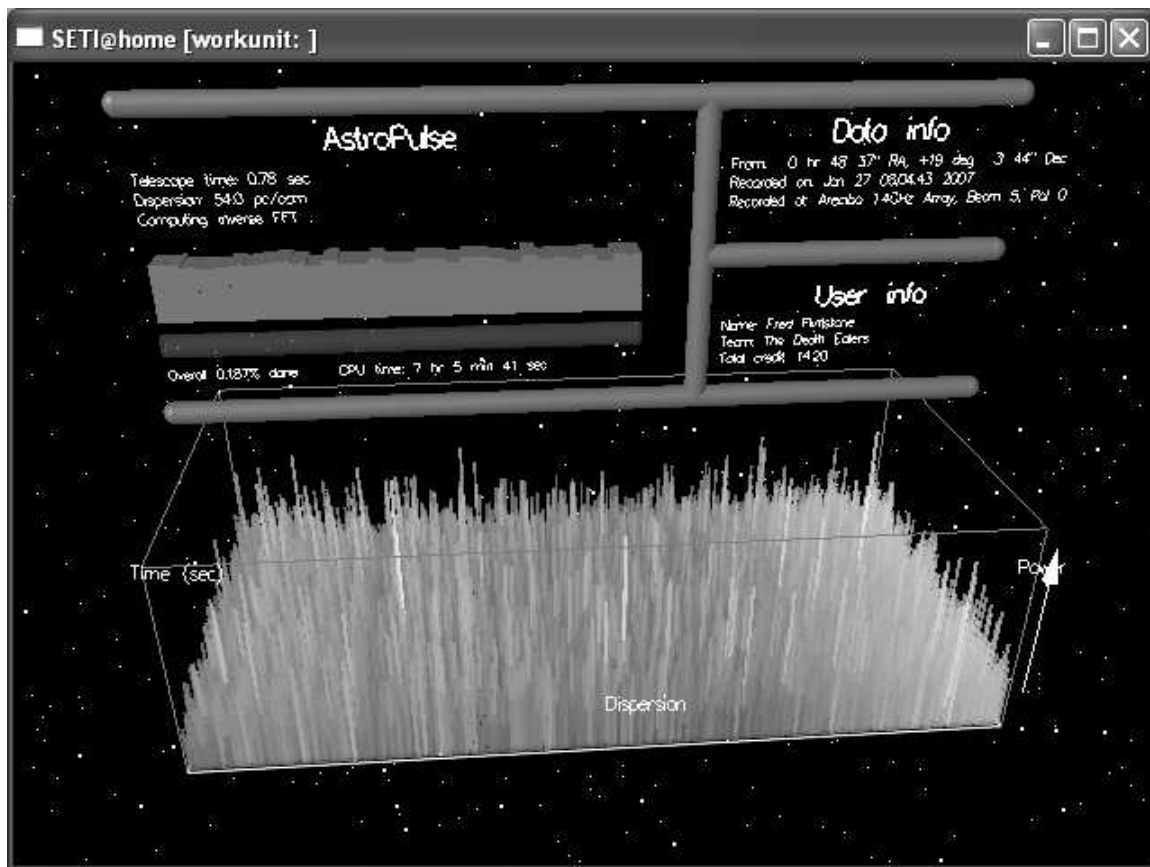


Fig. 2.— The Astropulse screen saver.